arXiv:1703.09906v1 [stat.ME] 29 Mar 2017

Modular Bayes screening for high-dimensional predictors

Yuhan Chen
Department of Statistical Science
Duke University, Durham, NC 27708
email: yuhan.chen@duke.edu

David Dunson
Department of Statistical Science
Duke University, Durham, NC 27708
email: dunson@duke.edu

March 30, 2017

Abstract

With the routine collection of massive-dimensional predictors in many application areas, screening methods that rapidly identify a small subset of promising predictors have become commonplace. We propose a new MOdular Bayes Screening (MOBS) approach, which involves several novel characteristics that can potentially lead to improved performance. MOBS first applies a Bayesian mixture model to the marginal distribution of the response, obtaining posterior samples of mixture weights, cluster-specific parameters, and cluster allocations for each subject. Hypothesis tests are then introduced, corresponding to whether or not to include a given predictor, with posterior probabilities for each hypothesis available analytically conditionally on unknowns sampled in the first stage and tuning parameters controlling borrowing of information across tests. By marginalizing over the first stage posterior samples, we avoid under-estimation of uncertainty typical of two-stage methods. We greatly simplify the model specification and reduce computational complexity by using *modularization*. We provide basic theoretical support for this approach, and illustrate excellent performance relative to competitors in simulation studies and the ability to capture complex shifts beyond simple differences in means. The method is illustrated with applications to genomics by using a very high-dimensional cis-eQTL dataset with roughly 38 million SNPs.

KEYWORDS: Genomics; High-dimensional; Independent screening; Large p small n; Mixture model; Modularization; Nonparametric Bayes; Variable selection.

1. INTRODUCTION

In modern scientific research, it has become routine to collect massive dimensional data for each study subject, leading to a massive p small to moderate n statistical problem. In most scientific studies collecting high-dimensional data, the over-arching interest is not in developing a black-box model for prediction; instead the focus is on variable selection. In particular, the scientists would like to select the subset of features out of a high-dimensional set of p candidates x_{ij}, which are predictive of a response y_i, with $i = 1, \ldots, n$ indexing the subject under study. There is a vast recent literature on high-dimensional feature/variable selection, with the overwhelming focus on one of two strategies. The first is to use a penalized optimization method, such as Lasso (Tibshirani 1996), SCAD (Fan and Li 2001), Dantzig selector (Candes and Tao 2007) or other variants, to obtain a sparse estimate of the regression coefficient vector. In general, such an approach has adequate performance only when the feature matrix X satisfies a number of stringent conditions, which seldom hold in scientific applications even approximately. When these assumptions are violated, performance can degrade badly.

This fact has lead to an emphasis on the second strategy, which is to first apply an independent screening algorithm using measures of association between x_{ij} and y_i separately for $j = 1, \ldots, p$, and then select features based on thresholding the corresponding p-values, test statistics or association estimates to maintain a superset containing most of the important features along with a number of false positives. This strategy was proposed by Fan and Lv (2008) using marginal correlation for sure independence screening (SIS). Fan and Song (2010) extended the idea to generalized linear models and Fan et al. (2011) considered nonparametric additive models. Zhu et al. (2011) introduced model-free screening via sure independent ranking and screening (SIRS). Other model-free procedures include the distance correlation screening (DCS) in Li et al. (2012) and the fused Kolmogorov filter in Mai and Zou (2013) and Mai and Zou (2015).

The focus of this article is on improving upon high-dimensional screening methodology through using Bayesian nonparametric and hierarchical models to avoid parametric assumptions on the data generating model, while borrowing information across the many tests being conducted. A potentially key disadvantage of most existing screening methods is the lack of any borrowing of information across the different tests. A novel idea we propose, which should have impact beyond

the screening case, is to place a nonparametric Bayes model on the *marginal* distribution of the response, with inclusion tests for each predictor then defined conditionally on unknowns in the marginal distribution. This approach has the dual advantages of avoiding parametric assumptions on the distribution of the response, while also automatically borrowing information across the screening tests through their shared conditioning on a common set of unknowns. In addition to this novel type of borrowing, we additionally place a more convention Bayesian hierarchical structure on the probability of variable inclusion, related to Scott and Berger (2010). A key advantage of the proposed formulation is that we can characterize uncertainty in posterior computation for the marginal response distribution via integrating conditional posterior probabilities of variable inclusion across first stage posterior samples. This is computationally scalable and avoids the under-estimation of uncertainty typical of two-stage procedures..

One caveat about our proposed approach is that it is not a coherent fully Bayes probability model in that we ignore certain dependencies for simplicity, robustness and computational tractability. In particular, to define a fully Bayes model for the marginal distribution of the response y in settings involving predictors $x_j = (x_{1j}, \ldots, x_{nj})'$, we would require a model for the conditional distribution of the response y given all the predictors x along with a model for the joint distribution of the predictors x. In very large p settings, computational and complexity considerations force a focus on relatively simple settings involving strong constraints (e.g., Gaussian, linear, highly sparse, etc). In most settings, such constraints are inconsistent with available prior information, and results are therefore highly questionable. Instead, we take the more mild approach of ignoring information on x in defining the marginal distribution of y. This is a type of *modularization* (Liu et al. 2009). The idea of modularization is that Bayesian models can be defined in *modules*, with posterior computation in certain modules not taking into account the model structure and data in other modules. In our case, the model for the marginal density $f(y)$ is one module, and we do not attempt to take into account information about the predictors in fitting this module. We suspect that closely related ideas may be a game changer in making (approximately) Bayesian approaches practical in modern massive dimensional data settings.

The existing literature on nonparametric Bayes variable selection is sparse. There are several methods for two group comparisons (Dunson and Peddada 2008; Ma and Wong 2011; Holmes et al.

2015), but such methods are computationally intensive for each group, and would need to be applied separately p times. There are also several approaches that are designed for variable selection in Bayesian nonparametric models for conditional response distributions, $f(y|x)$. For example, Chung and Dunson (2009) proposed a stochastic search variable selection algorithm under a probit stick-breaking mixture model for $f(y|x)$. Kessler et al. (2014) and Yang and Dunson (2016) used tensor factorizations to characterize the conditional distribution in order to capture complex interactions between the predictors. Jiang et al. (2017) introduced a nonparametric test using a nonparametric Bayesian slice inversed regression based on modeling the conditional distribution of a covariate given the discretized response. Their method is less computationally intensive than previous attempts, but is restricted to univariate y. Finally, Filippi et al. (2016) tested for pairwise dependence between two variables using Dirichlet process mixtures, improving efficiency by running MCMC for each marginal in parallel.

In Section 2 we propose the general MOdular Bayes Screening (MOBS) framework, provide basic theoretical support, and sketch a general approach to computation. In Section 3 we focus on the case in which the response is continuous and univariate, motivating a location-scale mixture of Gaussians for the marginal distribution. Section 4 contains a simulation study assessing operating characteristics relative to a variety of high-dimensional screening methods. Section 5 applies the screening method to a massive-dimensional cis-eQTL dataset with approximately 38 million SNPs. Section 6 contains a discussion. Technical details are included in an Appendix.

2. MODULAR BAYES SCREENING FRAMEWORK

2.1 General framework

MOdular Bayes Screening (MOBS) uses a two stage set-up and starts with a mixture model for the marginal density $f(y)$ of the response. For subjects $i = 1, \ldots, n$, let y_i denote the response, which can be either univariate or multivariate. Let $x_i = (x_{i1}, \ldots, x_{ip})'$ denote a high-dimensional vector of categorical predictors for subject i, with $x_{ij} \in \{0, 1, \ldots, d_j - 1\}$ having d_j levels. Suppose that marginally we have $y_i \sim f$ independently for $i = 1, \ldots, n$; we refer to f as the *baseline* density. Treating this baseline density as unknown, we apply a finite over-fitted mixture model (Rousseau

and Mengersen 2011):

$$f(y) = \sum_{h=1}^{k} \omega_h \mathcal{K}(y|\theta_h), \tag{1}$$

$$(\omega_1, \cdots, \omega_k)' \sim \text{Dir}\left(\frac{\alpha}{k}, \ldots, \frac{\alpha}{k}\right), \quad \theta_h \sim P_0,$$

where $\mathcal{K}$ is the kernel density, $\omega = (\omega_1, \cdots, \omega_k)'$ are the weights, $\theta = (\theta_1, \cdots, \theta_k)'$ the kernel parameters, P_0 is the prior distribution for each θ_h and k is a conservative upper bound on the number of mixture components needed to produce an accurate approximation of the unknown density. Let $c = (c_1, \ldots, c_n)'$ denote a latent variable vector with $c_i \in \{1, \ldots, k\}$ being the cluster membership of y_i and $\Pr(c_i = h) = \omega_h$. As k increases, (1) converges weakly to a Dirichlet process mixture model. This represents a common approximation to the Dirichlet Process mixture model (Ishwaran and Zarepour 2002) facilitating practical implementation.

To generalize the baseline density to allow dependence on the jth predictor x_j, we let

$$f(y|x_j) = \sum_{h=1}^{k} \omega_{jh}(x_j) \mathcal{K}(y|\theta_{jh}(x_j)), \tag{2}$$

where $\omega_j(l) = (\omega_{j1}(l), \cdots, \omega_{jk}(l))'$ and $\theta_j(l) = (\theta_{j1}(l), \cdots, \theta_{jk}(l))'$, $l \in \{0, 1, \ldots, d_j - 1\}$, represent the weights and component parameters for subjects with $x_{ij} = l$. Expression (2) modifies the baseline density (1) to allow the weights and kernel parameters to vary with x_j. Let

$$H_{0j} : \omega_j(0) = \ldots = \omega_j(d-1) = \omega, \quad \theta_j(0) = \ldots = \theta_j(d-1) = \theta, \tag{3}$$

$$H_{1j} : \text{Not all } \omega_j(l) \text{ are the same and/or not all } \theta_j(l) \text{ are the same}$$

correspond to discarding and including the jth predictor, respectively. The alternative hypothesis H_{1j} is composed of three cases: H_{1j1}, where ω varies but θ remains the same, H_{1j2}, where θ varies but ω is the same and H_{1j3}, where both vary. By testing both the weights and the cluster parameters, MOBS is able to capture deeper and more complex shifts in distribution than traditional methods that focus primarily on the mean. Clearly, $H_{1j} = H_{1j1} \cup H_{1j2} \cup H_{1j3}$. Let $\kappa = (\kappa_0, \kappa_{11}, \kappa_{12}, \kappa_{13})'$

be the prior probabilities of the four subhypotheses, where κ_0 is the prior probability of H_{0j} and $\kappa_{11}, \kappa_{12}, \kappa_{13}$ are the prior probabilities of the alternative hypotheses $H_{1j1}, H_{1j2}, H_{1j3}$ respectively. We suppose that the weight and component parameter vectors specific to each level of each predictor are distributed randomly about their baseline values according to the hierarchical model:

$$\omega_j(l) \sim \text{Dir}(\tau_\omega \omega), \quad \theta_{jh}(l) \sim P_1, \tag{4}$$

where P_1 is a distribution centered on θ with scale parameter τ_θ and τ_ω is the Dirichlet precision controlling variability about the mean ω. Using a Dirichlet prior for $\omega_j(l)$ maintains conjugacy with the Multinomial distribution of the component memberships c and facilitates marginalizing out $\omega_j(l)$ when calculating the likelihood. For the same reason, P_1 should ideally be conjugate to the kernel density $\mathcal{K}$. The precisions $\tau = (\tau_\omega, \tau_\theta)$ control how much $f(y|x_j)$ changes with predictor x_j for each j such that H_{1j} holds. The hierarchical structure favors borrowing of information in learning $f(y|x_j)$ across different predictors and levels of x_j.

2.2 Algorithmic details

Using (3) allows direct calculation of an analytic form for $\Pr(H_{0j}|y, x_j, \Lambda, \Phi)$, where Φ denotes ω, c, θ, the baseline weights and component parameters, and Λ denotes $\kappa, \tau_\omega, \tau_\theta$, the hyperparameters controlling borrowing of information and multiplicity adjustment in hypothesis testing. The key idea in our proposed MOBS approach is to run MCMC for posterior computation only for the baseline nonparametric model for the marginal response density $f(y)$. We do not know the exact kernel memberships, weights or parameters, and there does not exist a tractable, analytic form for $\Pr(H_{0j}|y, x_j, \Lambda)$, so MOBS instead takes the Monte Carlo integration of $\Pr(H_{0j}|y, x_j, \Lambda, \Phi)$ over the samples of Φ generated from the posterior distribution under the baseline model.

Let $\text{L}(c, y|H_{1j}, x_j, \omega, \theta, \Lambda)$ and $\text{L}(c, y|H_{0j}, x_j, \omega, \theta, \Lambda)$ be the likelihood conditional on the baseline MCMC output and hyperparameters controlling borrowing of information and multiplicity adjustment in hypothesis testing for H_{1j} and H_{0j}, respectively. Then, we have

$$\text{L}(c, y|H_{1j}, x_j, \omega, \theta, \Lambda) = \sum_{t=1}^{3} \frac{\kappa_{1t}}{1 - \kappa_0} \text{L}(c, y|H_{1jt}, x_j, \omega, \theta, \Lambda),$$

where $\mathrm{L}(c, y|H_{1jt}, x_j, \omega, \theta, \Lambda)$ is the conditional likelihood for H_{1jt} for $t = 1, 2, 3$. The posterior probabilities of H_{0j} and H_{1jt} conditional on the baseline weights and parameters are then, respectively,

$$\begin{aligned}
\Pr(H_{0j}|\kappa, y, x_j, \Phi) &= \frac{1}{1 + \frac{\kappa_{11}}{\kappa_0}\mathrm{BF}_{11}(j) + \frac{\kappa_{12}}{\kappa_0}\mathrm{BF}_{12}(j) + \frac{\kappa_{13}}{\kappa_0}\mathrm{BF}_{13}(j)}, \quad (5)\\
\Pr(H_{1jt}|\kappa, y, x_j, \Phi) &= \frac{\kappa_{1t}}{\kappa_0}\mathrm{BF}_{1t}(j)\Pr(H_{0j}|\kappa, y, x_j, \Phi), \quad (6)
\end{aligned}$$

where for $t = 1, 2, 3$,

$$\mathrm{BF}_{1t}(j) = \frac{\mathrm{L}(c, y|H_{1jt}, x_j, \omega, \theta, \Lambda)}{\mathrm{L}(c, y|H_{0j}, x_j, \omega, \theta, \Lambda)}$$

is the Bayes factor in favor of H_{1jt} over H_{0j} conditional on the hyperparameters and the baseline weights and parameters.

We calculate the conditional likelihoods given the baseline unknowns and hyperparameters as follows. First, under H_{0j}, we have the simple form:

$$\begin{aligned}
\mathrm{L}(c, y|H_{0j}, x_j, \omega, \theta, \Lambda) &= \mathrm{L}(c|H_{0j}, x_j, \omega, \theta, \Lambda)\mathrm{L}(y|H_{0j}, x_j, \Phi, \Lambda)\\
&= \prod_{h=1}^{k} \omega_h^{n_{jh}} \prod_{i:c_i=h} \mathcal{K}(y_i|\theta_h),
\end{aligned}$$

where $n_j(l) = (n_{j1}(l), \ldots, n_{jk}(l))'$, with $n_{jh}(l)$ the number of subjects having $x_{ij} = l$ and belonging to cluster h, and $n_j = (n_{j1}, \ldots, n_{jk})' = \sum_{l=0}^{d_j-1} n_j(l)$ the total number of subjects allocated to each component. For the three alternative subhypotheses

$$\begin{aligned}
\mathrm{L}(c, y|H_{1j1}, x_j, \omega, \theta, \Lambda) &= \mathrm{L}(c|H_{1j1}, x_j, \omega, \theta, \Lambda)\mathrm{L}(y|H_{1j1}, x_j, \Phi, \Lambda)\\
&= \left[\prod_{l=0}^{d_j-1} \int \left\{\prod_{h=1}^{k} \omega_{jh}(l)^{n_{jh}(l)}\right\} \mathrm{Dir}(\omega_j(l)|\tau_\omega\omega) d\omega_j(l)\right] \prod_{h=1}^{k} \prod_{i:c_i=h} \mathcal{K}(y_i|\theta_h),\\
\mathrm{L}(c, y|H_{1j2}, x_j, \omega, \theta, \Lambda) &= \mathrm{L}(c|H_{1j2}, x_j, \omega, \theta, \Lambda)\mathrm{L}(y|H_{1j2}, x_j, \Phi, \Lambda)\\
&= \prod_{l=0}^{d_j-1} \left\{\prod_{h=1}^{k} \omega_h^{n_{jh}}\right\} \int \prod_{h=1}^{k} P_1(\theta_{jh}(l)) \prod_{i:c_i=h} \mathrm{L}(y_i|\Phi, \Lambda, \theta_j(l)) d\theta_{jh}(l),\\
\mathrm{L}(c, y|H_{1j3}, x_j, \omega, \theta, \Lambda) &= \mathrm{L}(c|H_{1j3}, x_j, \omega, \theta, \Lambda)\mathrm{L}(y|H_{1j3}, x_j, \Phi, \Lambda)
\end{aligned}$$

$$= \left[\prod_{l=0}^{d_j-1}\int\left\{\prod_{h=1}^{k}\omega_{jh}(l)^{n_{jh}(l)}\right\}\mathrm{Dir}(\omega_j(l)|\tau_\omega\omega)d\omega_j(l)\right]$$
$$\prod_{l=0}^{d_j-1}\int\prod_{h=1}^{k}P_1(\theta_{jh}(l))\prod_{i:c_i=h}\mathrm{L}(y_i|\Phi,\Lambda,\theta_j(l))d\theta_{jh}(l).$$

These results imply after some calculations that

$$\begin{aligned}
\mathrm{BF}_{11}(j) &= \frac{\prod_{l=0}^{d_j-1}\int\left\{\prod_{h=1}^{k}\omega_{jh}(l)^{n_{jh}(l)}\right\}\mathrm{Dir}(\omega_j(l)|\tau_\omega\omega)d\omega_j(l)}{\omega_1^{n_1}\omega_2^{n_2}\ldots\omega_k^{n_k}},\\
\mathrm{BF}_{12}(j) &= \frac{\prod_{l=0}^{d_j-1}\int\prod_{h=1}^{k}P_1(\theta_{jh}(l))\prod_{i:c_i=h}\mathrm{L}(y_i|\Phi,\Lambda,\theta_j(l))d\theta_{jh}(l)}{\prod_{h=1}^{k}\prod_{i:c_i=h}\mathcal{K}(y_i|\theta_h)},\\
\mathrm{BF}_{13}(j) &= \frac{\prod_{l=0}^{d_j-1}\int\prod_{h=1}^{k}P_1(\theta_{jh}(l))\prod_{i:c_i=h}\mathrm{L}(y_i|\Phi,\Lambda,\theta_j(l))d\theta_{jh}(l)}{\prod_{h=1}^{k}\prod_{i:c_i=h}\mathcal{K}(y_i|\theta_h)}\\
&\quad\frac{\prod_{l=0}^{d_j-1}\int\left\{\prod_{h=1}^{k}\omega_{jh}(l)^{n_{jh}(l)}\right\}\mathrm{Dir}(\omega_j(l)|\tau_\omega\omega)d\omega_j(l)}{\omega_1^{n_1}\omega_2^{n_2}\ldots\omega_k^{n_k}}\\
&= \mathrm{BF}_{11}(j)\mathrm{BF}_{12}(j).
\end{aligned}$$

Here, $\mathrm{BF}_{11}(j)$ provides a weight of evidence of changes in mixture weights with x_j, while $\mathrm{BF}_{12}(j)$ provides the same for different kernels. $\mathrm{BF}_{13}(j)$ provides evidence for both weights and kernels being different and interestingly ends up being the product of $\mathrm{BF}_{11}(j)$ and $\mathrm{BF}_{12}(j)$.

Algorithm 1 Modular Bayes Screening

1: **Inputs:**
   data X, y
2: **Initialize:**
   $\kappa_0, \kappa_{11}, \kappa_{12}, \kappa_{13}, \tau_\mu, \tau_\theta$
3: Obtain samples for Φ from the posterior distribution under baseline mixture model using MCMC or an alternative algorithm.
4: **repeat**
5: Estimate $\Pr(H_{0j}|y, x_j, \Lambda)$ and $\Pr(H_{1jt}|y, x_j, \Lambda)$ for $t = 1, 2, 3$ by averaging $\Pr(H_{0j}|y, x_j, \Lambda, \Phi)$ and $\Pr(H_{1jt}|y, x_j, \Lambda, \Phi)$, respectively, over the baseline Φ samples for each j.
6: Set $\kappa_0, \kappa_{11}, \kappa_{12}, \kappa_{13}$ using the average of the $\Pr(H_{0j}|\kappa, y, x_j, \Lambda)$, $\Pr(H_{1j1}|y, x_j, \Lambda)$, $\Pr(H_{1j2}|y, x_j, \Lambda)$, and $\Pr(H_{1j3}|y, x_j, \Lambda)$ estimated above, respectively.
7: **until** convergence

We initialize $\kappa_0 = 0.5$, $\kappa_{11} = \kappa_{12} = \kappa_{13} = 0.5/3$. Estimating κ involves iteratively taking the

average of the posterior probabilities for each subhypothesis, which is effectively an empirical Bayes procedure. Our proposed MOBS approach leverages on the simplicity of posterior computation for $f(y)$ and the simple analytic forms shown above via Algorithm 1.

3. UNIVARIATE CONTINUOUS RESPONSE

3.1 Computation

While the proposed method allows for complex, arbitrary responses, this section provides an example under a simple univariate, continuous y setting. The baseline density $f(y)$ is represented as a mixture of univariate location-scale Gaussian kernels:

$$f(y) \quad = \quad \sum_{h=1}^{k} \omega_h \mathrm{N}(y|\mu_h, \sigma_h^2), \tag{7}$$

where $\{\mu_h, \sigma_h^2\} = \theta_h$ are the mean and variance parameters of the component h with prior $P_0 = \mathrm{N}(\mu_h|\mu_0, q\sigma_h^2)\mathrm{IGa}(\sigma_h^2|a, b)$. Details on the Gibbs sampler for the base distribution can be found in the Appendix. The conditional density $f(y|x_j)$ can similarly be represented as

$$f(y|x_j) \quad = \quad \sum_{h=1}^{k} \omega_{jh}(x_j) \mathrm{N}(y|\mu_{jh}(x_j), \sigma_{jh}^2(x_j)), \tag{8}$$

where $\{\mu_{jh}(x_j), \sigma_{jh}^2(x_j)\} = \theta_{jh}(x_j)$ with prior $P_1 = \mathrm{N}(\mu_{jh}(l)|\mu_h, \sigma_{jh}^2(l)/\tau_\mu)\mathrm{IGa}(\sigma_{jh}^2(l)|a_h, b_h)$, with $a_h = \frac{\tau_\sigma}{\sigma_h^4}$, $b_h = \frac{\tau_\sigma}{\sigma_h^2}$ and the scale parameters $\tau_\theta = \{\tau_\mu, \tau_\sigma\}$.

To derive $\mathrm{BF}_{11}(j)$, note that for each j and l,

$$\int \left\{ \prod_{h=1}^{k} \omega_{jh}(l)^{n_{jh}(l)} \right\} \mathrm{Dir}(\omega_j(l)|\tau_\omega \omega) d\omega_j(l) = \frac{\beta(n_j(l) + \tau_\omega \omega)}{\beta(\tau_\omega \omega)},$$

where β is the multivariate beta function. Hence

$$\mathrm{BF}_{11}(j) = \frac{1}{\omega_1^{n_1} \omega_2^{n_2} \ldots \omega_k^{n_k}} \prod_{l=0}^{d_j - 1} \frac{\beta(n_j(l) + \tau_\omega \omega)}{\beta(\tau_\omega \omega)}. \tag{9}$$

Similarly to find $\mathrm{BF}_{12}(j)$ for each j, first note that

$$\prod_{h=1}^{k}\prod_{i:c_i=h}\mathcal{K}(y_i|\theta_h)=\frac{1}{(2\pi)^{\frac{n}{2}}}\prod_{h=1}^{k}\frac{1}{\sigma_h^{n_h}}e^{-\frac{\sum_{i:c_i=h}(y_i-\mu_h)^2}{2\sigma_h^2}}$$

and

$$\begin{aligned}
&\prod_{l=0}^{d_j-1}\int\prod_{h=1}^{k}P_1(\theta_{jh}(l))\prod_{i:c_i=h}\mathrm{L}(y_i|\Phi,\Lambda,\theta_j(l))d\theta_{jh}(l)\\
=&\prod_{l=0}^{d_j-1}\int\mathrm{L}(y|c,\omega,\mu_j(l),\sigma_j^2(l))p(\sigma_{jh}^2(l)|a_h,b_h)p(\mu_{jh}(l)|\mu_h,\tau_\mu,\sigma_{jh}^2(l))d\mu_j(l)d\sigma_j^2(l)\\
=&\frac{1}{(2\pi)^{\frac{n}{2}}}\prod_{l=0}^{d_j-1}\prod_{h=1}^{K}\frac{\Gamma(a_{jh}(l))}{\Gamma(a_h)}\frac{b_h^{a_h}}{b_{jh}(l)^{a_{jh}(l)}}\frac{\sqrt{\tau_\mu}}{\sqrt{\tau_\mu+n_{jh}(l)}},
\end{aligned}$$

where

$$\begin{aligned}
&a_{jh}(l)=a_h+\frac{n_{jh}(l)}{2},\quad \bar{y}_{jh}(l)=\frac{1}{n_{jh}(l)}\sum_{i:c_i=h,x_{ij}=l}y_i,\\
&b_{jh}(l)=b_h+\frac{n_{jh}(l)\tau_\mu}{2(\tau_\mu+n_{jh}(l))}\{\mu_h-\bar{y}_{jh}(l)\}^2+\sum_{i:c_i=h,x_{ij}=l}\frac{\{y_i-\bar{y}_{jh}(l)\}^2}{2}.
\end{aligned}$$

It follows that

$$\mathrm{BF}_{12}(j)=\left(\prod_{l=0}^{d_j-1}\prod_{h=1}^{k}\frac{\Gamma(a_{jh}(l))}{\Gamma(a_h)}\frac{b_h^{a_h}}{b_{jh}(l)^{a_{jh}(l)}}\frac{\sqrt{\tau_\mu}}{\sqrt{\tau_\mu+n_{jh}(l)}}\right)\left(\prod_{h=1}^{k}\sigma_h^{n_h}e^{\frac{\sum_{c_i=h}(y_i-\mu_h)^2}{2\sigma_h^2}}\right). \quad (10)$$

Running MOBS requires first generating samples of c,ω,μ and σ^2 from the posterior under the baseline mixture model. One then iteratively alternates between estimation of $\Pr(H_{0j}|y,x_j,\Lambda)$ using (9) and (10) over the samples of the baseline c,ω,μ and σ^2 and estimation of κ as the average of the estimated posterior probabilities. The strength of MOBS lies in the ease of computing $\mathrm{BF}_{11}(j)$ and $\mathrm{BF}_{12}(j)$ in a trivially parallelizable manner. The overall algorithm has computational complexity $O(np)$. The approach can be trivially modified to accommodate multivariate and discrete settings; in such cases, expression (10) will take a different form.

3.2 Hyperparameters selection

An important question when running MOBS is the selection of hyperparameters, especially that of the precision parameters $\tau = \{\tau_\mu, \tau_\sigma, \tau_\omega\}$. The various τ can be interpreted as the precisions controlling the distance between the density of y under the null hypothesis and that of $y|x_j = l$ under the alternative hypothesis. Naturally, higher values of τ would suggest a smaller distance, as $y|x_j = l$ converges to y when $\tau_\mu, \tau_\sigma, \tau_\omega \to \infty$.

We aim to tune τ such that the prior signal-to-noise ratio is within a reasonable range somewhere between 0.05 and 0.1. For any fixed $j = 1, 2, ..., p$ and $l = 0, 1, ..., d_j - 1$, let $\tilde{\Phi} = \{\mu, \sigma, \omega, \mu_j(l), \sigma_j(l), \omega_j(l)\}$. Given $\tilde{\Phi}$, the densities of μ, y and $y|x_j = l$ are

$$\begin{aligned} f_{\mu|\tilde{\Phi}}(t) &= \sum_{h=1}^{k} \omega_h \mathrm{N}(t|\mu_0, q\sigma_h^2), \\ f_{y|\tilde{\Phi}}(t) &= \sum_{h=1}^{k} \omega_h \mathrm{N}(t|\mu_h, \sigma_h^2), \\ f_{y|x_j=l,\tilde{\Phi}}(t) &= \sum_{h=1}^{k} \omega_{jh}(l) \mathrm{N}(t|\mu_{jh}(l), \sigma_{jh}^2(l)). \end{aligned}$$

Let

$$\Delta_0 = E\left[\|f_{y|\tilde{\Phi}} - f_{\mu|\tilde{\Phi}}\|_2^2\right],$$

$$\Delta_1 = E\left[\|f_{y|\tilde{\Phi}} - f_{y|x_j=l,\tilde{\Phi}}\|_2^2\right]$$

be the average square of L_2-norm distances between the densities of $\mu|\tilde{\Phi}$ and $y|\tilde{\Phi}$ and between those of $y|\tilde{\Phi}$ and $y|x_j = l, \tilde{\Phi}$ respectively. Since $\int \mathrm{N}(t|\mu_1, \sigma_1^2)\mathrm{N}(t|\mu_2, \sigma_2^2)dt = N\left(\mu_1 - \mu_2|0, \sigma_1^2 + \sigma_2^2\right)$, we know

$$\begin{aligned} \|f_{y|\tilde{\Phi}} - f_{\mu|\tilde{\Phi}}\|_2^2 &= \sum_{h=1}^{k}\sum_{t=1}^{k} \omega_h\omega_t \left[\mathrm{N}(0|0, q(\sigma_h^2 + \sigma_t^2)) + \mathrm{N}(\mu_h - \mu_t|0, \sigma_h^2 + \sigma_t^2)\right. \\ &\quad \left. -2\mathrm{N}(\mu_h - \mu_0|0, \sigma_h^2 + q\sigma_t^2)\right], \end{aligned}$$

$$\begin{aligned}\|f_{y|\tilde{\Phi}} - f_{y|x_j=l,\tilde{\Phi}}\|_2^2 &= \sum_{h=1}^{k}\sum_{t=1}^{k}\omega_h\omega_t \mathrm{N}(\mu_h - \mu_t|0, \sigma_h^2 + \sigma_t^2) \\ &\quad + \sum_{h=1}^{k}\sum_{t=1}^{k}\omega_{jh}(l)\omega_{jt}(l)\mathrm{N}(\mu_{jh}(l) - \mu_{jt}(l)|0, \sigma_{jh}^2(l) + \sigma_{jt}^2(l)) \\ &\quad -2\sum_{h=1}^{k}\sum_{t=1}^{k}\omega_h\omega_{jt}(l)\mathrm{N}(\mu_h - \mu_{jt}(l)|0, \sigma_h^2 + \sigma_{jt}(l)^2).\end{aligned}$$

While Δ_0 and Δ_1 are both intractable, we can estimate their values by averaging $\|f_{y|\tilde{\Phi}} - f_{\mu|\tilde{\Phi}}\|_2^2$ and $\|f_{y|\tilde{\Phi}} - f_{y|x_j=l,\tilde{\Phi}}\|_2^2$ over samples of $\tilde{\Phi}$. Empirically, we find that for a given k, setting $\mu_0 = 0, \alpha = k, a = 2, b = 0.02, q = 50, \tau_\mu = 50, \tau_\sigma = 50, \tau_\omega = k^{3/2} + 8(k-1)$ provides a reasonable default specification that is relatively stable and has a signal-to-noise ratio Δ_1/Δ_0 between 0.05 and 0.1.

3.3 Theoretical properties

We initially study asymptotic properties of MOBS treating parameters Φ as known, and then consider sure screening consistency under general settings. We first investigate properties of $\Pr(H_{0j}|y, x_j, \Lambda, \Phi)$. The posterior odds conditioned on Φ and Λ are

$$\frac{\Pr(H_{1j}|y, x_j, \Lambda, \Phi)}{\Pr(H_{0j}|y, x_j, \Lambda, \Phi)} = \frac{\kappa_{11}}{\kappa_0}\mathrm{BF}_{11}(j) + \frac{\kappa_{12}}{\kappa_0}\mathrm{BF}_{12}(j) + \frac{\kappa_{13}}{\kappa_0}\mathrm{BF}_{12}(j)\mathrm{BF}_{11}(j).$$

Set $d_j = 2$. For $l = 0, 1$, $h = 1, 2, \ldots, k$ and $j = 1, 2, \ldots, p$, let

$$n_j^l = \sum_{h=1}^{k} n_{jh}(l), \qquad \lambda_0 = \frac{n_j^0}{n_j^0 + n_j^1}, \qquad p_{jhl} = \frac{n_{jh}(l)}{n_j^l}.$$

Therefore, $n_j^0 = n\lambda_0$ and $n_j^1 = n(1 - \lambda_0)$ where $n = n_j^0 + n_j^1$. Let $\Omega_j(l)$ and $\Theta_j(l) = (M_j(l), \Sigma_j^2(l))$ be the true values of $\omega_j(l)$ and $\theta_j(l) = (\mu_j(l), \sigma_j^2(l))$, respectively, for $j = 1, 2, \ldots, p, l = 0, 1$. Then under H_{0j}, $\Omega_j(l) = \omega$, $M_j(l) = \mu$ and $\Sigma_j^2(l) = \sigma^2$.

Theorem 3.1 *Let $n_j^0, n_j^1 \to \infty$, with $\frac{n_j^0}{n_j^1} = \frac{\lambda_0}{1-\lambda_0}$ for some fixed λ_0. Treat Φ as known.*

(1) When H_{0j} is true, we have $BF_{11}(j) \to 0$ with rate $n^{-(k-1)}$ and $BF_{12}(j) \to 0$ with rate n^{-2k}. These two estimates imply that $Pr(H_{1j}|y, x_j, \Lambda, \Phi) \to 0$ with rate $n^{-(k-1)}$.

(2) Under the condition $\Omega_j(0) \neq \Omega_j(1)$, we have $\log BF_{11}(j) \to \infty$ with rate An and under the

condition $\Theta_j(0) \neq \Theta_j(1)$*, we have* $\log BF_{12}(j) \to \infty$ *with rate* Bn *where*

$$\begin{aligned} A &= \sum_{h=1} \left[\lambda_0 \Omega_{jh}(0) \log \left\{ \frac{\Omega_{jh}(0)}{\omega_h} \right\} + \lambda_1 \Omega_{jh}(1) \log \left\{ \frac{\Omega_{jh}(1)}{\omega_h} \right\} \right] \\ B &= \frac{1}{2} \sum_{h=1}^{k} \sum_{l=0}^{1} \frac{\lambda_l \Omega_{jh}(l) \{M_{jh}(l) - \mu_h\}^2}{\sigma_h^2} + \frac{1}{2} \sum_{h=1}^{k} \sum_{l=0}^{1} \lambda_l \Omega_{jh}(l) \left[\frac{\Sigma_{jh}^2(l)}{\sigma_h^2} - 1 + \log \left\{ \frac{\sigma_h^2}{\Sigma_{jh}^2(l)} \right\} \right]. \end{aligned}$$

Therefore under H_{1j}*,* $Pr(H_{0j}|y, x_j, \Lambda, \Phi) \to 0$ *with rate* $e^{-n \max(A,B)}$*.*

The above theorem shows rates of convergence of hypothesis probabilities and Bayes factors for the different hypotheses in the event that the true weights, parameters and cluster allocations are known. However, in practice MOBS uses posterior samples for the baseline parameters to account for uncertainty. Now consider the situation where the parameters are unknown and the model may be misspecified. In particular, we extend our above results using misspecification techniques given in Kleijn and van der Vaart (2006) and used in Lock and Dunson (2015). Let $\mathbb{F}$ be the set of all convex combinations of Gaussian distributions $\{f_h(\cdot|\theta_h)\}$ where $\theta_h = \{\mu_h, \sigma_h^2\}$ and let P define a prior on $\mathbb{F}$. Let f_0 be the true distribution and f^* the closest convex combination in $\mathbb{F}$ to f_0 under Kullback-Leibler divergence. We define $B(\epsilon, f^*|f_0)$ to be a neighborhood of the density f^* under the measure induced by the density f_0:

$$B(\epsilon, f^*|f_0) = \left\{ f \in \mathbb{F} : -\int f_0 \log \frac{f}{f^*} \leq \epsilon^2, \int f_0 \log \left(\frac{f}{f^*} \right)^2 \leq \epsilon^2 \right\}$$

and define $d(f_1, f_2)$ to be the weighted Hellinger distance

$$d^2(f_1, f_2) = \frac{1}{2} \int (f_1^{1/2} - f_2^{1/2})^2 \frac{f_0}{f^*}.$$

Lemma 3.2 *Let* $y_1, \ldots, y_n$ *be independent with density* f_0*. Assume* $f^* = argmin_{f \in \mathbb{F}} KL(f_0||f)$ *exists and* $P_n(B(\epsilon, f^*|f_0)) > 0$ *for all* $\epsilon > 0$ *where* P_n *is the posterior measure given prior* P*. Let* $\omega^* = (\omega_1^*, \ldots, \omega_k^*)'$ *be the component weights and* $\theta^* = (\theta_1^*, \ldots, \theta_k^*)'$ *be the kernel parameters corresponding to* f^* *and let* ω *and* θ *be the true weights and kernel parameters. Assume* ω^* *and* θ^* *are unique in that* $\sum \omega_h f_h(\cdot|\theta_h) = \sum \omega_h^* f_h(\cdot|\theta_h^*) = f^*$ *only if* $\omega = \omega^*, \theta = \theta^*$*. Then for any fixed*

$\epsilon > 0$,

$$P_n((\omega, \theta) \in (\mathbb{S}^{k-1}, \mathbb{R}^k) : ||(\omega, \theta) - (\omega^*, \theta^*)|| \geq \epsilon | y) \to 0.$$

Theorem 3.3 *For $l = \{0, 1\}$, let $y_i | x_{ij} = l$ denote the subset of y_i with size n_j^l such that $x_{ij} = l$. Assume $y_i | x_{ij} = 0$ are independent with density $f^{(0)}$ and $y_i | x_{ij} = 1$ are independent with density $f^{(1)}$. Let*

$$f^{*(0)} = argmin_{f \in \mathbb{F}} KL(f^{(0)} || f), \quad f^{*(1)} = argmin_{f \in \mathbb{F}} KL(f^{(1)} || f).$$

Assume the uniqueness condition for the previous theorem holds for $f^{*(0)}$ and $f^{*(1)}$. If $f^{(0)} = f^{(1)}$, $Pr(H_{0j} | y, x_j, \Lambda) \to 1$ as $n_j^0, n_j^1 \to \infty$ and if $f^{*(0)} \neq f^{*(1)}$, $Pr(H_{0j} | y, x_j, \Lambda) \to 0$ as $n_j^0, n_j^1 \to \infty$.

This result suggests that our posterior probability is consistent under H_{0j} and holds under weak conditions under H_{1j} for each $j = 1, \ldots, p$. However, consistency can fail under H_{1j} when $f^{(0)}$ and $f^{(1)}$ have the same closest point f^* under K-L divergence in $\mathbb{F}$. In practice, $f^{(0)}$ and $f^{(1)}$ would need to be extremely close to have the same closest point, so this is a very mild condition.

Next, we use the previous theorem to establish sure screening consistency for MOBS. Following standard screening conventions (Li et al. 2012), let $F(y|X)$ denote the conditional distribution function of y given the predictors X. Define

$$D = \{j : F(y|X) \text{ depends on } x_j \text{ for some } y\}$$

to be the true set of relevant predictors. A good screening method can identify a small subset S such that $D \subset S$. Existing literature have focused predominantly on frequentist methods that select S based on rankings or thresholding using some test statistic such as marginal correlation in SIS or distance correlation in DCS. MOBS instead relies on thresholding of the estimated posterior null probabilities $\hat{\pi}_j = \Pr(H_{0j} | y, x_j, \Lambda)$ for $j = 1, \ldots, p$. In particular, let

$$\hat{D} = \{j : \hat{\pi}_j \leq \hat{\pi}_{(d_n)}\},$$

where $\hat{\pi}_{(1)}, \ldots, \hat{\pi}_{(p)}$ are the order values and d_n is an integer. Let $\gamma_j = 1$ if the jth predictor is marginally related to y and $\gamma_j = 0$ otherwise. We impose the following two conditions:

(C1) All jointly important predictors $j \in D$ are also marginally important (i.e. $\gamma_j = 1$) and $d_n \geq |D|$.

(C2) All jointly important predictors $j \in D$ satisfy the constraint that $f_j^{(0)}$ and $f_j^{(1)}$ are not the closest in K-L divergence to the same $f_j^* \in \mathbb{F}$.

Condition 1 is similar to conditions found in the screening literature (Fan and Lv 2008), restricting metrics of marginal importance for important predictors to be non-zero. Condition 2 requires Theorem 3.3 to be satisfied, and is mild as noted above.

Theorem 3.4 *Under conditions (C1) and (C2), when* $n, p \to \infty$,

$$Pr(D \subset \hat{D}) \to 1.$$

4. SIMULATION

4.1 Screening accuracy

In this section, we assess the performance of MOBS against five existing screening methods using simulated datasets. The first competitor is sure independence screening or SIS (Fan and Lv 2008), and we use the `SIS R` package. We also compare against three frequentist model-free methods, sure independent ranking and screening or SIRS (Zhu et al. 2011), using code found at *http://users.stat.umn.edu/ wangx346/research/example1b.txt*, the distance correlation screening or DCS (Li et al. 2012), using the `energy R` package and the fused Kolmogorov filter (Mai and Zou 2015) using code provided by the authors. Finally, we compare to the JYL Bayesian nonparametric test (Jiang et al. 2017) with code found at *http://www.people.fas.harvard.edu/ junliu/BF/bfslice.html*.

We simulate 100 replicates for all methods under six different models. Under these settings, we test $n = 200$ observations and $p = 2000$ total predictors and generate $x_{ij} \in \{0, 1\}$ and y. First, we consider the case where X is independent and generate the features with $x_{ij} \sim \text{Bern}(0.5)$, $i = 1, \ldots, n$, $j = 1, \ldots, p$ for all 2000 predictors. However, in biological applications (e.g. involving SNPs), X can often exhibit moderate correlations in blocks of predictors. To mimic such settings, we randomly select 600 predictors and draw them in the following fashion with correlation $\rho = 0.5$. We first draw $z = (z_{ij})_{n \times p}$ for each $z_{ij} \sim N(0, 1)$ for $i = 1, \ldots, n$ and $j = 1, 2, \ldots, p$. Next, construct

$b = (b_{ij})_{n \times p}$ by replacing the first $\rho \cdot n$ rows of z such that for $1 \leq j \leq p$ for some constant a,

$$b_{ij} = \begin{cases} z_{i1} + a z_{ij}, & 1 \leq i \leq \rho \cdot n \\ z_{ij}, & \rho \cdot n < j \leq n \end{cases}$$

To convert this continuous data to multinomial, for each predictor, we can simply assign them by quantiles. The remaining 1900 predictors are generated from independent Gaussians without correlation.

Model 1 (*Linear regression*, *uncorrelated*) : $y = 1+2x_1+x_2-2x_3+x_4-2x_5+\epsilon$, where $\epsilon \sim \mathrm{N}(0,1)$ is independent of X and X is uncorrelated.

Model 2 (*Linear regression*, *correlated*) : y is generated the same as with *Model 1* but with a block of predictors with correlation described above.

Model 3 (*Single index*, *uncorrelated*) : $y = (1+2x_1+x_2-2x_3+x_4-2x_5)^2+\epsilon$, where $\epsilon \sim \mathrm{N}(0,1)$ is independent of X and X is uncorrelated.

Model 4 (*Single index*, *correlated*) : y is generated the same as with *Model 3* but with a block of predictors with correlation described above.

Model 5 (*Mixture*, *uncorrelated*) : y is generated using 6 predictors at positions chosen at random from $\{1, \ldots, p\}$. Let S be the true vector of jointly important indices. The resulting marginal density is then sampled from 64 mixtures of Gaussian distributions with mean $\mu \sim \mathrm{Unif}(-1, 1)$ and standard deviation $\sigma \sim \mathrm{Unif}(0, 1/8)$ each corresponding to different permutations of S. If $x_{iS} = (x_{iS1}, \ldots, x_{iS6})$, we subsequently draw $y_i \sim N(\mu_{x_{iS1},\ldots,x_{iS6}}, \sigma^2_{x_{iS1},\ldots,x_{iS6}})$. X is generated to be uncorrelated.

Model 6 (*Mixture*, *correlated*) : y is generated the same as with *Model 5* but with a block of predictors with correlation described above.

In order to implement MOBS, we use the computational strategy described in Section 2 and Section 3 and set the deafult hyperparameters $\mu_0 = 0, q = 50, a = 2, b = 0.02, \alpha = k, \tau_\omega = k^{3/2} + 8(k-1), \tau_\mu = \tau_\sigma = 50$, with $k = 3$ for the linear case since we expect the response to be mostly unimodal and $k = 7$ for the more multimodal single index and mixture models. We run the Markov chain of the baseline model 6000 times and perform Monte Carlo integration using the last 500 samples. If potential issues involving label-switching for the samples of θ arise, we use methods

found in Stephens (2000) and Papastamoulis (2016).

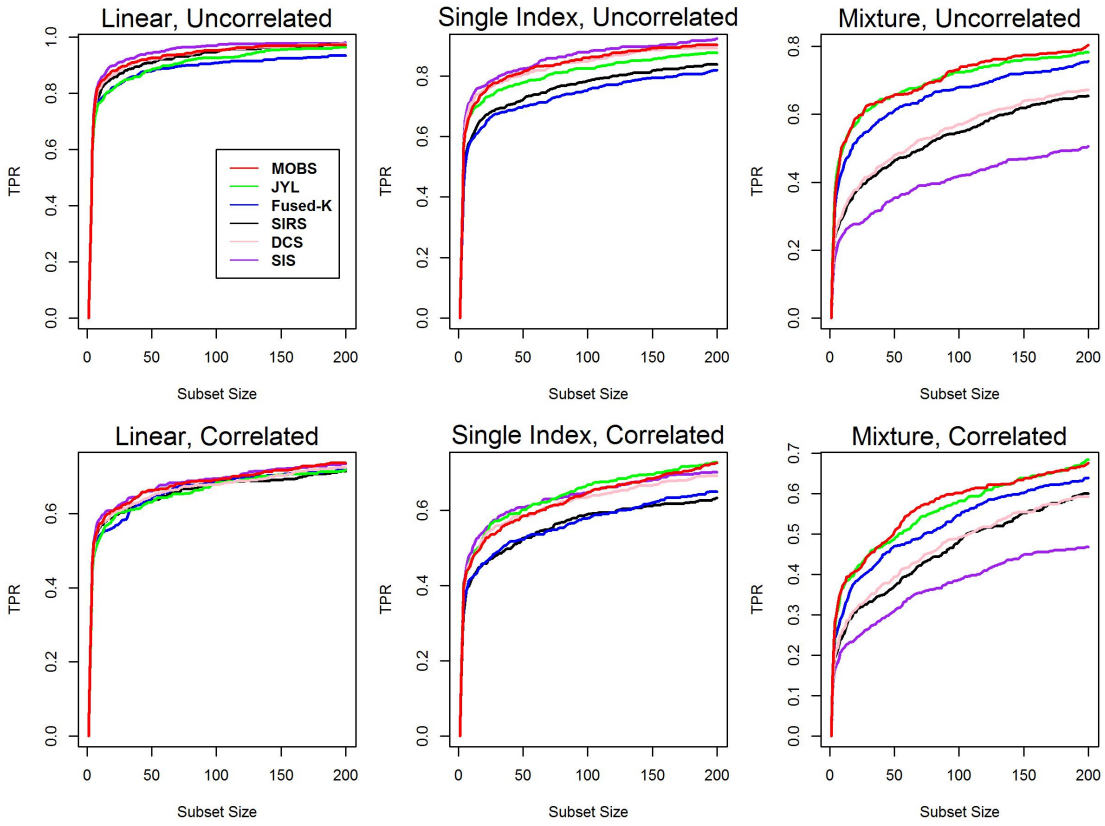

Figure 1: ROC curves for all six methods under linear, single index and mixture settings, both correlated and uncorrelated.

Figure 1 summarizes the ROC curves for the six different methods by comparing the true positive rate against the false positive rate. In each of the simulations, MOBS exhibits the best or one of the best performance in comparison to the SIS, SIRS, DCS, the fused Kolmogorov filter and the JYL approaches. In the mixture setting, the nonparametric methods dominate as expected. For the linear regression case and single index case, the other two nonparametric methods, the fused Kolmogorov filter and JYL both demonstrate weakness in these settings, but MOBS still performs strongly and is comparable to SIS. In a multimodal mixture setting, both ω and θ play large roles in detecting change, whereas in a unimodal setting, ω is less relevant and MOBS reduces to essentially a t-test.

4.2 Computational efficiency

Understanding the computational efficiency of variable screening is important as the various methods should be scalable to huge data sets. We measure the total time taken to run each of the six methods under two different settings using data generated from an uncorrelated linear regression model setting in *Model 1* but with varying n and p. Our first example fixes the number of subjects at 250 and then alter the number of predictors. Next, we fix the number of predictors at 2500 and instead change the number of subjects.

Figure 2 summarize the results under both methods. As expected, the linear predictor SIS exhibits significantly faster performance as a simple parametric model though MOBS remains competitive against the other nonparametric and model-free screening methods. All six screening methods are simple to parallelize and scale linearly with the number of predictors p. However, in regards to the number of samples n, MOBS scales linearly unlike slower competitors such as JYL and DCS.

5. APPLICATIONS

5.1 Data

In this section, we illustrate our approach by analyzing the GEUVADIS cis-eQTL dataset (Lappalainen et al. 2013), publicly available at *http://www.ebi.ac.uk/Tools/geuvadis-das/*. The data set consists of messenger RNA and microRNA on lymphoblastoid cell line (LCL) samples from 462 individuals provided by the 1000 Genomes Project along with roughly 38 million SNPs. These individuals are taken from 5 different populations: the Yoruba (YRI), CEPH (CEU), Toscani (TSI), British (GBR) and Finns (FIN). We focus on the gene $E2F2$ (Ensemble ID: ENSG00000007968) as our response. $E2F2$ plays a key role in the control of the cell cycle (Attwooll et al. 2004). Traditional eQTL analysis is often limited by the daunting size of the dataset, so our objective is to use screening to select a subset of SNPs that are associated with the gene.

5.2 Results

We first ignore the SNPs. We standardize the response $E2F2$ and then run the baseline model under the univariate continuous setting using 7000 iterations with 6500 burn-in and default hyperparameters described in the previous sections with $k = 5$ clusters. Figure 3 contains a histogram

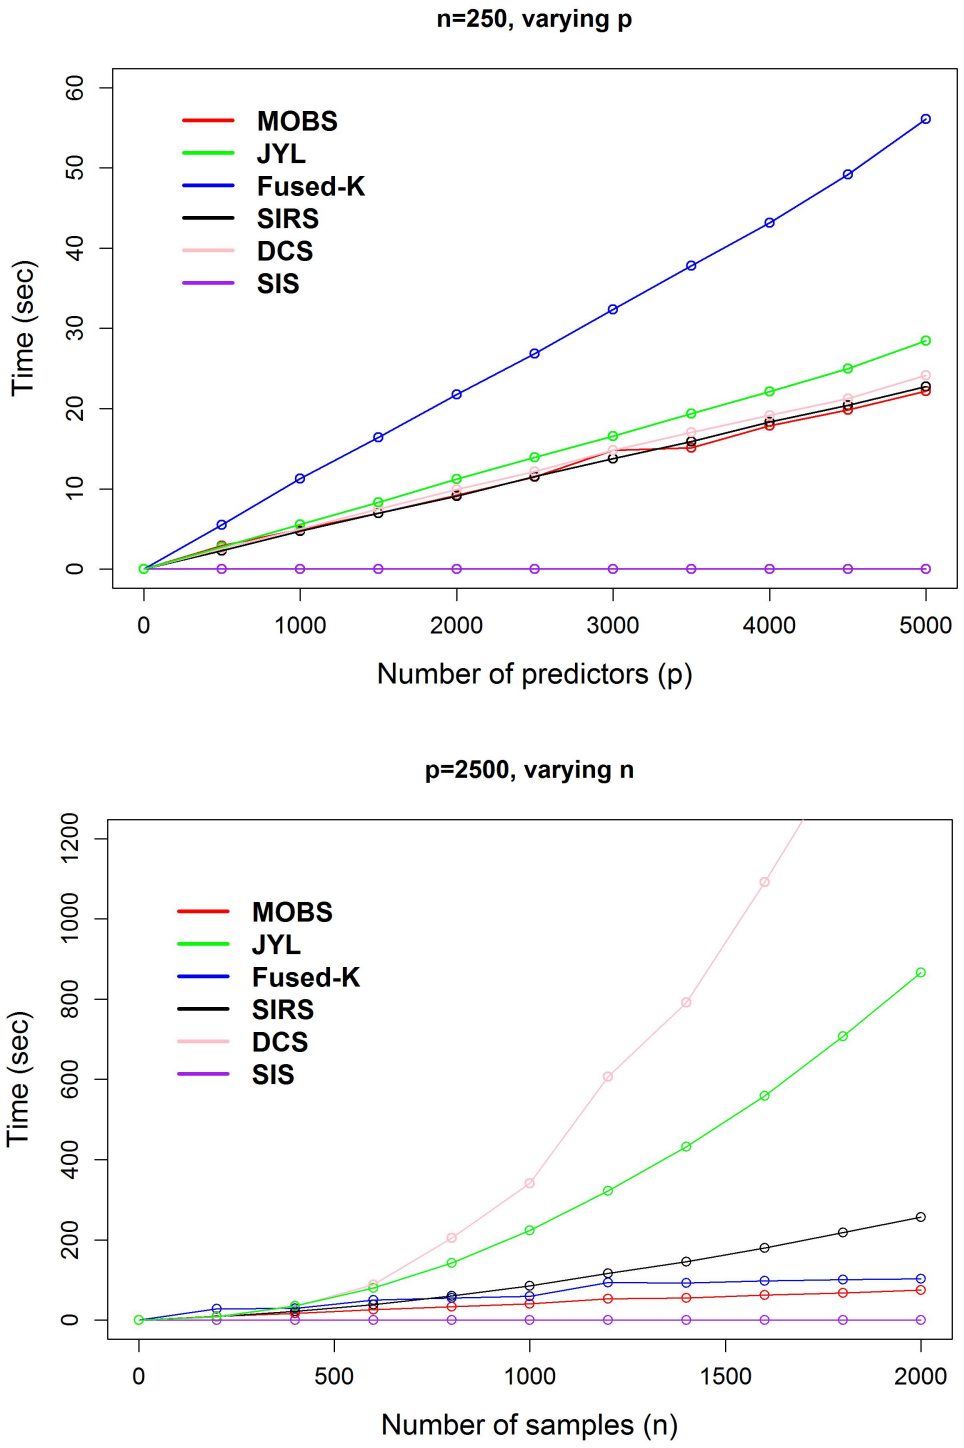

Figure 2: Computational speed against number of samples n when $p = 2500$ and computational speed against number of predictors p when $n = 250$.

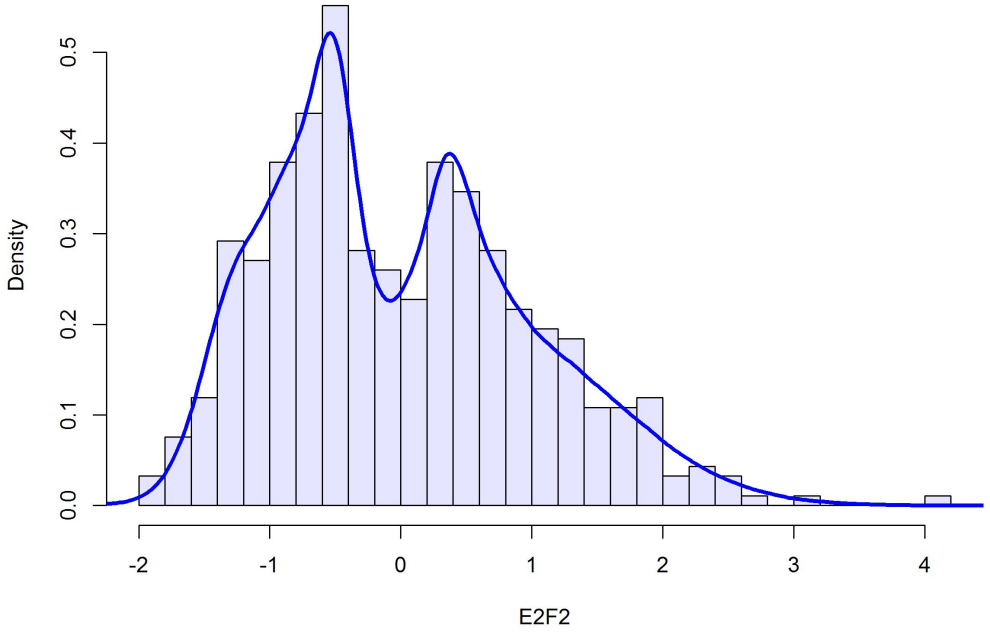

Figure 3: Histogram of $E2F2$ (Ensemble ID: ENSG00000007968) with the estimated density from sampling the baseline model with $k = 5$ clusters.

of the standardized $E2F2$ as well as the estimated marginal density from the Gibbs sampler in the first stage.

Next, we run the second stage and estimate $\Pr(H_{0j}|y, x_j, \Lambda, \Phi)$ by averaging $\Pr(H_{0j}|y, x_j, \Lambda, \Phi)$ over the baseline Φ samples, after first removing the SNPs that contain missing data or have x_j values being all 0, all 1 or all 2. The resulting posterior probabilities are displayed in Figure 4. Note that the majority of posterior probabilities are concentrated near 1, with roughly 90% having a posterior probability of 0.95 or greater. On the other hand, about 0.4% of posterior probabilities have values 0.05 or less, with the smallest being $1.864872 \cdot 10^{-8}$. Figure 5 shows the estimated densities at four selected SNPs where red represents samples with $x_j = 0$, green with $x_j = 1$ and blue with $x_j = 2$. The SNPs are chosen with their $\Pr(H_{0j}|y, x_j, \Lambda)$ at four different values of under-0.0001, 0.36, 0.61 and over-0.99. Overall, the results suggest that MOBS performs well. SNPs with $\Pr(H_{0j}|y, x_j, \Lambda)$ close to 1 have all three estimated distributions nearly identical consistent with the raw data histograms. On the other hand, SNPs with $\Pr(H_{0j}|y, x_j, \Lambda)$ close to 0 have three

highly varying estimated densities, which are again consistent with smoothed versions of the data histograms.

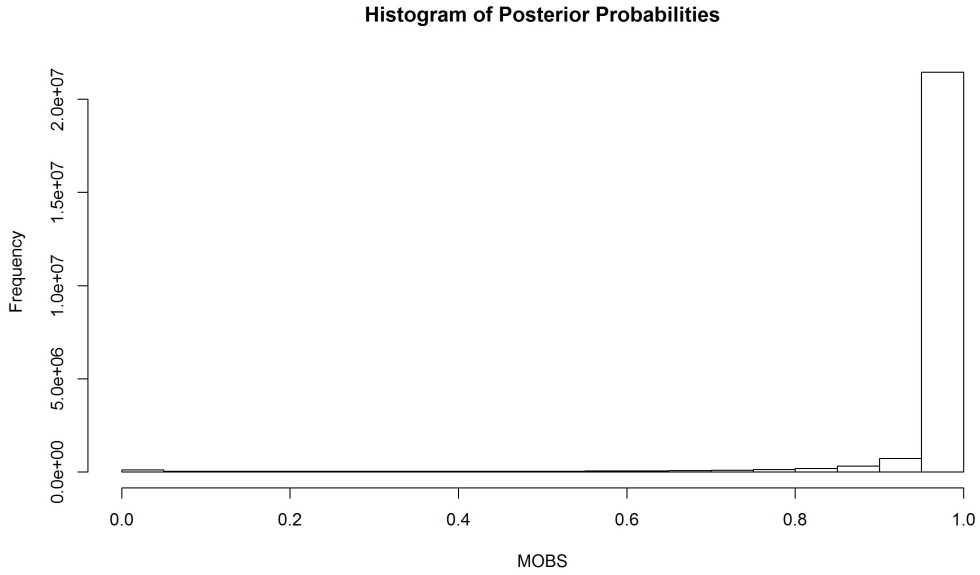

Figure 4: Histogram of $\Pr(H_{0j}|X, y)$, measuring the association between $E2F2$ and the SNPs.

Traditional screening methods generally focus on testing the differences in mean. However, such methods miss more complex shifts in variance or density shape. MOBS is sensitive to not only mean shifts but also shifts in the shape and variance. To judge how much of the SNP selection is due to differences in means, we run the six methods MOBS, DCS, SIS, SIRS, the fused Kolmogorov filter and the JYL Bayesian nonparametric test and then select the 50 most significant predictors across all six models. For DCS, SIS, SIRS, the fused Kolmogorov filter, we convert each x_j into $x_j^{(0)}$ and $x_j^{(1)}$ where $x_j^{(0)}$ and $x_j^{(1)}$ are 0/1 indicators for $x_j = 0$ and $x_j = 1$ respectively. We then run the models on both $x_j^{(0)}$ and $x_j^{(1)}$ and take the overall test statistic x_j as the maximum of the two. Across all models, we take the sum of the absolute value of the mean for each of the three groups $x_j = 0, 1$ and 2 and display their values in Table 1. The results show that the SNPs selected by DCS, SIS, SIRS and JYL have higher values of total absolute mean distance, which imply that they are predominantly identifying the shifts in the mean. On the other hand, MOBS and FUSEDK have significantly lower values and hence are focusing not just on the mean but on more complex differences.

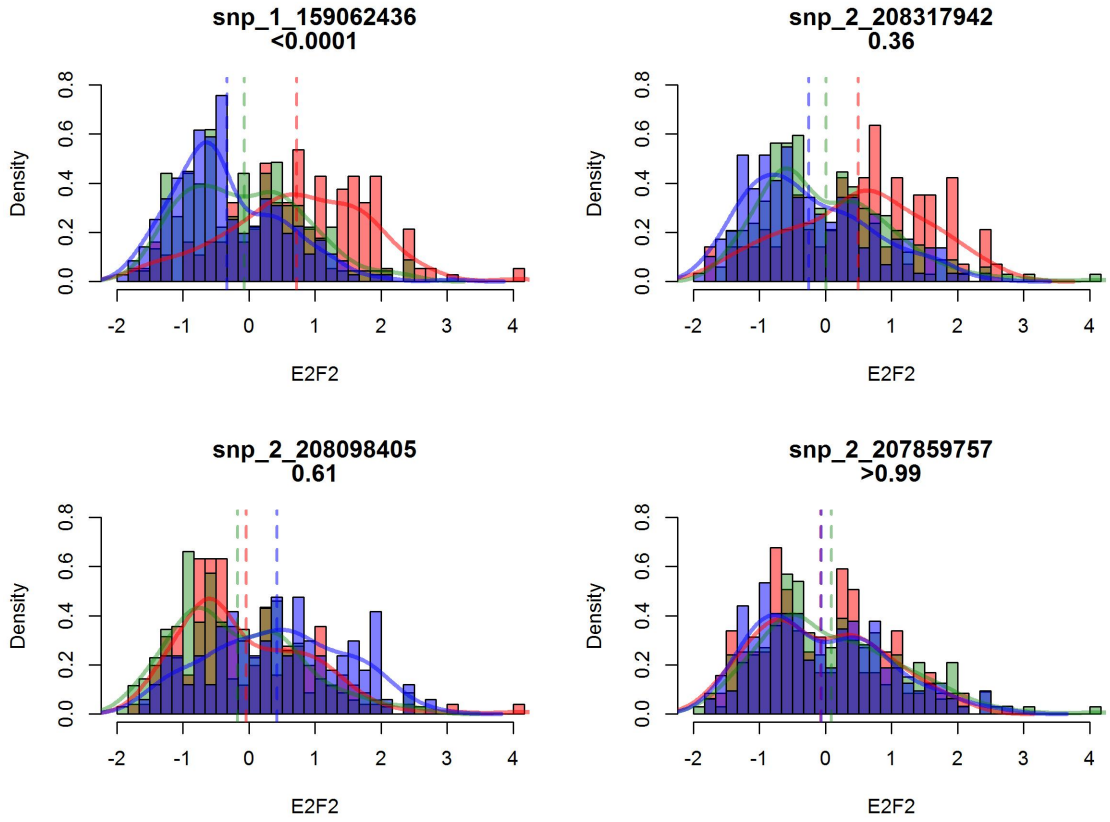

Figure 5: The estimated densities for $E2F2$ for $x_j = 0$ (red), $x_j = 1$ (green) and $x_j = 2$ (blue) at four selected SNPs with varying values of $\Pr(H_{0j}|x_j, y, \Lambda)$ at under 0.0001, 0.36, 0.61 and over 0.99. Histograms are also displayed in red, green and blue for each of the three groups.

Finally, we assess whether the SNPs identified by MOBS have good predictive performance compared with those identified by DCS, SIS, SIRS, FUSEDK and JYL. We split the data with 370 samples for training and 92 for testing and run 5-fold cross validation. For each screening method, we select the 50 most significant SNPs. In order to examine the predictive capabilities of each method, we fit LASSO after screening. Therefore, the resulting algorithms are MOBS-L, DCS-L, SIS-L, SIRS-L, FUSEDK-L and the JYL-L and their average measure squared errors (MSE)

	MOBS	DCS	SIS	SIRS	FUSEDK	JYL
Total Absolute Mean Distance	1.12	1.88	1.91	1.86	1.14	1.64

Table 1: The sum of the absolute value of the mean of each cluster for the six models.

are given in Table 2. While the nonparametric MOBS-L, JYL-L and FUSEDK-L perform better than the SIS-L, DCS-L and SIRS-L, all six methods suffer from poor predictive performance with relatively high MSE. This is not surprising as it tends to be very difficult to accurately predict gene expression based on SNPs alone. The main focus of the analysis is on identifying a promising set of SNPs for further study and not on prediction. A nice feature of MOBS is that it can pick up SNPs different from those identified by existing screening methods.

	MOBS-L	DCS-L	SIS-L	SIRS-L	FUSEDK-L	JYL-L
Average MSE	0.850	0.882	0.880	0.886	0.847	0.852

Table 2: Comparison of the predictive performance of the six models using the 50 most significant predictors over $E2F2$.

6. DISCUSSION

In this paper, we introduce MOBS, a Bayesian nonparametric class of screening procedures for arbitrary response y and categorical predictors $x_i = (x_{i1}, \ldots, x_{ip})'$. By using a type of modularization, MOBS first runs an MCMC algorithm over a baseline model for the marginal response density $f(y)$ and then uses these baseline samples to rapidly estimate each conditional density $f(y|x_j)$. Through simulations, MOBS performs very competitively in terms of both screening performance and computational cost. An analysis on the cis-eQTL genomics dataset reveals that MOBS is able to capture complex shifts beyond a simple difference in means, while maintaining a competitive predictive performance.

The main focus of the paper is on introducing our novel modular screening technique, and providing an illustration through univariate continuous outcome data with massive-dimensional categorical predictors; the cis-eQTL dataset analyzed in Section 5 involved roughly 38 million SNP predictors, which is orders of magnitude larger than is considered in most articles on "high-dimensional" methods. The proposed MOBS strategy can be applied much more broadly; for example, it is straightforward to consider outcome variables that are multivariate and complex, as long as a Bayesian mixture model and MCMC algorithm has been developed for data of that type. This is the case not just for multivariate and mixed scale data but also for object data ranging from shapes and curves to graphs. It is also possible to adapt the methods to accommodate continuous or mixed categorical and continuous predictors. One approach is define d pre-specified knots for

each continuous predictor (e.g., spaced at quantiles of the empirical distribution of x_j) and use kernels to interpolate $f(y|x_j)$ between these knots.

Acknowledgements

The authors thank Barbara Engelhardt of Princeton University for generously providing the cis-eQTL dataset used in Section 5. This research was partially funded by grant 1546130 of the US National Science Foundation (NSF) and grant 3130624 of the US Office for Naval Research (ONR).

REFERENCES

Attwooll, C., Denchi, E. L., and Helin, K. (2004), "The E2F Family: Specific Functions and Overlapping Interests," *The EMBO Journal*, 23, 4709–4716.

Candes, E. and Tao, T. (2007), "The Dantzig Selector: Statistical Estimation When p Is Much Larger Than n," *The Annals of Statistics*, 35, 2313–2351.

Chung, Y. and Dunson, D. B. (2009), "Nonparametric Bayes Conditional Distribution Modeling with Variable Selection," *Journal of the American Statistical Association*, 104, 1646–1660.

Dunson, D. B. and Peddada, S. D. (2008), "Bayesian Nonparametric Inference on Stochastic Ordering," *Biometrika*, 95, 859–874.

Fan, J. and Li, R. (2001), "Variable Selection via Nonconcave Penalized Likelihood and Its Oracle Properties," *Journal of the American Statistical Association*, 96, 1348–1360.

Fan, J. and Lv, J. (2008), "Sure Independence Screening for Ultrahigh Dimensional Feature Space (with Discussion)," *Journal of the Royal Statistical Society, Series B*, 70, 849–911.

Fan, J. and Song, R. (2010), "Sure Independence Screening in Generalized Linear Models with NP-Dimensionality," *The Annals of Statistics*, 38, 3567–3604.

Fan, J., Feng, Y., and Song, R. (2011), "Nonparametric Independence Screening in Sparse Ultra-High Dimensional Additive Models," *Journal of the American Statistical Association*, 106, 544–557.

Filippi, S., Holmes, C., and Nieto-Barajas, L. (2016), "Scalable Bayesian Nonparametric Measures for Exploring Pairwise Dependence via Dirichlet Process Mixtures," *Electronic Journal of Statistics*, 10, 3338–3354.

Holmes, C. C., Caron, F., Griffin, J. E., and Stephens, D. A. (2015), "Two-Sample Bayesian Nonparametric Hypothesis Testing," *Bayesian Analysis*, 10, 297–320.

Ishwaran, H. and Zarepour, M. (2002), "Exact and Approximate Sum Representations for Dirichlet Process," *The Canadian Journal of Statistics*, 30, 269–283.

Jiang, B., Ye, C., and Liu, J. S. (2017), "Bayesian Nonparametric Tests via Sliced Inverse Modeling," *Bayesian Analysis*, 12, 89–112.

Kessler, D. C., Taylor, J. A., and Dunson, D. B. (2014), "Learning Phenotype Densities Conditional on Many Interacting Predictors," *Bioinformatics*, 30, 1562–1568.

Kleijn, B. J. K. and van der Vaart, A. W. (2006), "Misspecification in Infinite-Dimensional Bayesian Statistics," *The Annals of Statistics*, 34, 837–877.

Lappalainen, T., Sammeth, M., Friedländer, M. R., t Hoen, P. A., Monlong, J., Rivas, M. A., Gonzàlez-Porta, M., Kurbatova, N., Griebel, T., Ferreira, P. G., Barann, M., Wieland, T., Greger, L., van Iterson, M., Almlöf, J., Ribeca, P., Pulyakhina, I., Esser, D., Giger, T., Tikhonov, A., Sultan, M., Bertier, G., MacArthur, D. G., Lek, M., Lizano, E., Buermans, H. P., Padioleau, I., Schwarzmayr, T., Karlberg, O., Ongen, H., Kilpinen, H., Beltran, S., Gut, M., Kahlem, K., Amstislavskiy, V., Stegle, O., Pirinen, M., Montgomery, S. B., Donnelly, P., McCarthy, M. I., Flicek, P., Strom, T. M., Lehrach, H., Schreiber, S., Sudbrak, R., Carracedo, Á., Antonarakis, S. E., Häsler, R., Syvänen, A. C., van Ommen, G. J., Brazma, A., Meitinger, T., Rosenstiel, P., Guigó, R., Gut, I. G., Estivill, X., and Dermitzakis, E. T. (2013), "Transcriptome and Genome Sequencing Uncovers Functional Variation in Humans," *Nature*, 501, 506–511.

Li, R., Zhong, W., and Zhu, L. (2012), "Feature Screening via Distance Correlation Learning," *Journal of the American Statistical Association*, 107, 1129–1139.

Liu, F., Bayarri, M. J., and Berger, J. (2009), "Modularization in Bayesian Analysis, with Emphasis on Analysis of Computer Models," *Bayesian Analysis*, 4, 119–150.

Lock, E. and Dunson, D. B. (2015), "Shared Kernel Bayesian Screening," *Biometrika*, 102, 829–842.

Ma, L. and Wong, W. H. (2011), "Coupling Optional Pólya Trees and the Two Sample Problem," *Journal of the American Statistical Association*, 106, 1553–1565.

Mai, Q. and Zou, H. (2013), "The Kolmogorov Filter for Variable Screening in High-Dimensional Binary Classification," *Biometrika*, 100, 229–234.

Mai, Q. and Zou, H. (2015), "The Fused Kolmogorov Filter: A Nonparametric Model-Free Screening Method," *The Annals of Statistics*, 43, 1471–1497.

Papastamoulis, P. (2016), "Label.Switching: An R Package for Dealing with the Label Switching Problem in MCMC Outputs," *Journal of Statistical Software*, 69, 1–24.

Rousseau, J. and Mengersen, K. (2011), "Asymptotic Behaviour of the Posterior Distribution in Overfitted Mixture Models," *Journal of the Royal Statistical Society, Series B*, 73, 689–710.

Scott, J. G. and Berger, J. O. (2010), "Bayes and Empirical-Bayes Multiplicity Adjustment in the Variable-Selection Problem," *The Annals of Statistics*, 38, 2587–2619.

Stephens, M. (2000), "Dealing with Label Switching in Mixture Models," *Journal of the Royal Statistical Society, Series B*, 62, 795–809.

Tibshirani, R. (1996), "Regression Shrinkage and Selection via the Lasso," *Journal of the Royal Statistical Society, Series B*, 58, 267–288.

Yang, Y. and Dunson, D. B. (2016), "Bayesian Conditional Tensor Factorizations for High-Dimensional Classification," *Journal of the American Statistical Association*, 111, 656–669.

Zhu, L., Li, L., Li, R., and Zhu, L. (2011), "Model-Free Feature Screening for Ultrahigh-Dimensional Data," *Journal of the American Statistical Association*, 106, 1464–1475.

Appendix

Algorithm for univariate Gaussian base mixture model

In this section, we provide algorithmic details to sample from the base univariate Gaussian mixture model described in **Algorithm 1**. Let $n_h = \sum_{i=1}^{n} 1_{c_i=h}$ be the number of individuals in component h and $\bar{y}_h = \sum_{i:c_i=h} y_i/n_h$ be the mean of y within component h. We use a Gibbs sampler to update parameters using the full conditionals:

1. Sample the cluster allocations from multinomial conditional distributions with

$$\Pr(c_i = h|-) = \frac{\omega_h N(y_i|\mu_h, \sigma_h^2)}{\sum_{l=1}^{k} \omega_l N(y_i|\mu_l, \sigma_l^2)}, \quad h = 1, 2, \ldots, k$$

2. Sample the component-specific means and variances from

$$\mu_h, \sigma_h^2|- \sim \mathrm{N}(\mu_h|\hat{\mu}_h, \hat{q}_h\sigma_h^2)\mathrm{IGa}(\sigma_h^2|\hat{a}_h, \hat{b}_h), \quad h = 1, 2, \ldots, k$$

   where

$$\hat{q}_h = (q^{-1} + n_h)^{-1}, \quad \hat{\mu}_h = \hat{q}_h(q^{-1}\mu_0 + n_h\bar{y}_h), \quad \hat{a}_h = a + n_h/2,$$

$$\hat{b}_h = b + \frac{1}{2}\left[\sum_{i:c_i=h} (y_i - \bar{y}_h)^2 + \left(\frac{n_h}{1 + qn_h}\right)(\bar{y}_h - \mu_0)^2\right],$$

3. Sample the vector of weights on the different mixture components

$$\omega_1, \ldots, \omega_k|- \sim \mathrm{Dir}\left(\frac{\alpha}{k} + n_1, \ldots, \frac{\alpha}{k} + n_k\right).$$

Proof of Theorem 3.1

Lemma 6.1 *As $n_j^0, n_j^1 \to \infty$, with $\frac{n_j^0}{n_j^1} = \frac{\lambda_0}{1-\lambda_0}$ for some fixed λ_0, we have*

$$BF_{11}(j) \sim C_\omega(j) n^{-(k-1)} \prod_{h=1}^{k} \left(\frac{p_{jh0}}{\omega_h}\right)^{n_{jh}(0)} \left(\frac{p_{jh1}}{\omega_h}\right)^{n_{jh}(1)},$$

where

$$C_\omega(j) = \frac{(2\pi)^{k-1}}{\left(\beta(\tau_\omega\omega)\right)^2}\left(\lambda_0(1-\lambda_0)\right)^{-\frac{k-1}{2}}\prod_{h=1}^{k}\left(p_{jh0}p_{jh1}\right)^{\tau_\omega\omega_h-\frac{1}{2}},$$

and

$$BF_{12}(j) \sim C_\theta(j)e^{-n/2}n^{-2k}\prod_{h=1}^{k}\left(\frac{n_{jh}(0)}{2b_{jh}(0)}\right)^{a+\frac{n_{jh}(0)}{2}}\left(\frac{n_{jh}(1)}{2b_{jh}(1)}\right)^{a+\frac{n_{jh}(1)}{2}}(\sigma_h^2)^{\frac{n_{jh}(0)+n_{jh}(1)}{2}}e^{\sum_{i:c_i=h}\frac{(y_i-\mu_h)^2}{2\sigma_h^2}},$$

where

$$C_\theta(j) = (2\pi)^k 2^{2k}\left(\lambda_0(1-\lambda_0)\right)^{-k}\tau_\mu^k\prod_{h=1}^{k}\frac{{b_h}^{2a_h}}{\Gamma(a_h)^2 p_{jh0}p_{jh1}}.$$

Proof. By Stirling's formula, for $c>0$ and $x\to\infty$, we have

$$\begin{aligned}\Gamma(x) &\sim \sqrt{\frac{2\pi}{x}}\left(\frac{x}{e}\right)^x,\\ \Gamma(x+c) &\sim x^c\Gamma(x).\end{aligned}$$

Thus,

$$\Gamma(x+c) \sim \sqrt{2\pi}\,\frac{x^{x+c-\frac{1}{2}}}{e^x}.$$

When $n_j^0\to\infty$, we know

$$\begin{aligned}\beta(n_j(0)+\tau_\omega\omega) &= \frac{\prod_{h=1}^{k}\Gamma(n_{jh}(0)+\tau_\omega\omega_h)}{\Gamma(n_j^0+\tau_\omega)}\\ &\sim \frac{\prod_{h=1}^{k}\sqrt{2\pi}\frac{(n_{jh}(0))^{n_{jh}(0)+\tau_\omega\omega_h-\frac{1}{2}}}{e^{n_{jh}(0)}}}{\sqrt{2\pi}\frac{(n_j^0)^{n_j^0+\tau_\omega-\frac{1}{2}}}{e^{n_j^0}}}\\ &\sim \left(\frac{2\pi}{n_j^0}\right)^{\frac{k-1}{2}}\frac{\prod_{h=1}^{k}(n_{jh}(0))^{n_{jh}(0)+\tau_\omega\omega_h-1/2}}{(n_j^0)^{n_j^0+\tau_\omega-k/2}}\\ &= \left(\frac{2\pi}{n_j^0}\right)^{\frac{k-1}{2}}\prod_{h=1}^{k}p_{jh0}^{n_{jh}(0)+\tau_\omega\omega_h-1/2}\\ &= \left(\frac{2\pi}{n\lambda_0}\right)^{\frac{k-1}{2}}\prod_{h=1}^{k}p_{jh0}^{n_{jh}(0)+\tau_\omega\omega_h-1/2}\end{aligned}$$

Therefore,

$$\begin{aligned}
\mathrm{BF}_{11}(j) &= \frac{\beta(n_j(0)+\tau_\omega\omega)}{\beta(\tau_\omega\omega)}\frac{\beta(n_j(1)+\tau_\omega\omega)}{\beta(\tau_\omega\omega)}\frac{1}{\omega_1^{n_1},\ldots,\omega_k^{n_k}} \\
&\sim \left[\left(\frac{2\pi}{n\lambda_0}\right)^{\frac{k-1}{2}}\prod_{h=1}^{k} p_{jh0}^{n_{jh}(0)+\tau_\omega\omega_h-1/2}\right]\left[\left(\frac{2\pi}{n(1-\lambda_0)}\right)^{\frac{k-1}{2}}\prod_{h=1}^{k} p_{jh1}^{n_{jh}(1)+\tau_\omega\omega_h-1/2}\right] \\
&\quad \frac{1}{(\beta(\tau_\omega\omega))^2\omega_1^{n_1},\ldots,\omega_k^{n_k}} \\
&= C_\omega(j)n^{-(k-1)}\prod_{h=1}^{k}\left(\frac{p_{jh0}}{\omega_h}\right)^{n_{jh}(0)}\left(\frac{p_{jh1}}{\omega_h}\right)^{n_{jh}(1)},
\end{aligned}$$

where

$$C_\omega(j) = \frac{(2\pi)^{k-1}}{(\lambda_0(1-\lambda_0))^{\frac{k-1}{2}}(\beta(\tau_\omega\omega))^2}\prod_{h=1}^{k}(p_{jh0}p_{jh1})^{\tau_\omega\omega_h-1/2}.$$

As for $\mathrm{BF}_{12}(j)$, we see that

$$\Gamma(a_h+n_{jh}(0)/2) \sim \sqrt{2\pi}\ \frac{(n_{jh}(0)/2)^{a_h+n_{jh}(0)/2-1/2}}{e^{n_{jh}(0)/2}}.$$

Thus

$$\begin{aligned}
&\prod_{h=1}^{k}\Gamma(a_h+n_{jh}(0)/2)\Gamma(a_h+n_{jh}(1)/2) \\
&\sim (2\pi)^k e^{-n/2}\prod_{h=1}^{k}\left(\frac{n_{jh(0)}}{2}\right)^{a_h+n_{jh}(0)/2-1/2}\left(\frac{n_{jh}(1)}{2}\right)^{a_h+n_{jh}(1)/2-1/2}.
\end{aligned}$$

Also

$$\frac{\tau_\mu}{\sqrt{(\tau_\mu+n_{jh}(0))(\tau_\mu+n_{jh}(1))}} \sim \frac{\tau_\mu}{\sqrt{n_{jh}(0)n_{jh}(1)}}.$$

Thus

$$\begin{aligned}
\mathrm{BF}_{12}(j) &\sim \frac{\prod_{h=1}^{k}\frac{\Gamma(a_{jh}(0))\Gamma(a_{jh}(1))}{\Gamma(a_h)\Gamma(a_h)}\frac{{b_h}^{a_h}{b_h}^{a_h}}{{b_{jh}(0)}^{a_{jh}(0)}{b_{jh}(1)}^{a_{jh}(1)}}\frac{\tau_\mu}{\sqrt{n_{jh}(0)n_{jh}(1)}}}{\prod_{h=1}^{k}\frac{1}{(\sigma_h^2)^{(n_{jh}(0)+n_{jh}(1))/2}}e^{-\sum_{c_i=h}(y_i-\mu_h)^2/(2\sigma_h^2)}} \\
&\sim C_\theta(j)e^{-n/2}n^{-2k}\frac{\prod_{h=1}^{k}\left(\frac{n_{jh}(0)}{2b_{jh}(0)}\right)^{a+n_{jh}(0)/2}\left(\frac{n_{jh}(1)}{2b_{jh}(1)}\right)^{a+n_{jh}(1)/2}}{\prod_{h=1}^{k}\frac{1}{(\sigma_h^2)^{(n_{jh}(0)+n_{jh}(1))/2}}e^{-\sum_{c_i=h}(y_i-\mu_h)^2/(2\sigma_h^2)}}
\end{aligned}$$

where

$$C_\theta(j) = (2\pi)^k 2^{2k} (\lambda_0(1-\lambda_0))^{-k} \tau_\mu^k \prod_{h=1}^{k} \frac{{b_h}^{2a_h}}{\Gamma(a_h)^2 p_{jh0} p_{jh1}}.$$

□

To prove Theorem 3.1, we first set $\lambda_1 = 1 - \lambda_0$ and fix $j = 1, \ldots, p$.

We see that for each $h = 1, 2, .., k$ and $l = 0, 1$, for given $\Omega_{jh}(l)$, the integer $n_{jh}(l)$ is the number of $\{y_i\}$ with $x_{ij} = l$ in component h. As the event that $y_i | x_{ij} = l$ belongs to component h has probability $\Omega_{jh}(l)$, by Central Limit Theorem, we know

$$p_{jhl} - \Omega_{jh}(l) = O_p\left(\frac{1}{\sqrt{n_j^l}}\right). \tag{11}$$

We see that for small t,

$$(1+t)\log(1+t) = t + O(t^2). \tag{12}$$

By (11), (12), we can get

$$\frac{p_{jhl}}{\Omega_{jh}(l)} \log\left[\frac{p_{jhl}}{\Omega_{jh}(l)}\right] = \frac{p_{jhl}}{\Omega_{jh}(l)} - 1 + O_p\left(\frac{1}{n_j^l}\right) \tag{13}$$

Similarly,

$$\bar{y}_{jh}(l) - M_{jh}(l) = O_p\left(\frac{1}{\sqrt{n_{jh}(l)}}\right), \tag{14}$$

$$\frac{1}{n_{jh}(l)} \sum_{i:c_i=h, x_{ij}=l} \{y_i - \bar{y}_{jh}(l)\}^2 = \Sigma_{jh}^2(l) + O_p\left(\frac{1}{\sqrt{n_{jh}(l)}}\right). \tag{15}$$

Since

$$\sum_{i:c_i=h, x_{ij}=l} \{y_i - \bar{y}_{jh}(l)\}^2 \;=\; \sum_{i:c_i=h, x_{ij}=l} \{y_i - M_{jh}(l)\}^2 - n_{jh}(l)\{\bar{y}_{jh}(l) - M_{jh}(l)\}^2.$$

By (14 and (15),

$$\begin{aligned}
\frac{2b_{jh}(l)}{n_{jh}(l)} &= \frac{\sum_{i:c_i=h,x_{ij}=l}\{y_i-\bar{y}_{jh}(l)\}^2}{n_{jh}(l)}+O_p\left(\frac{1}{\sqrt{n_{jh}(l)}}\right)\\
&= \frac{\sum_{i:c_i=h,x_{ij}=l}\{y_i-M_{jh}(l)\}^2}{n_{jh}(l)}-\{\bar{y}_{jh}(l)-M_{jh}(l)\}^2+O_p\left(\frac{1}{\sqrt{n_{jh}(l)}}\right)\\
&= \Sigma^2_{jh}(l)+O_p\left(\frac{1}{\sqrt{n_{jh}(l)}}\right),
\end{aligned}$$

and

$$\begin{aligned}
\frac{n_{jh}(l)}{2}\log\left[\frac{2b_{jh}(l)}{n_{jh}(l)\Sigma^2_{jh}(l)}\right] &= \frac{n_{jh}(l)}{2}\left[\frac{2b_{jh}(l)}{n_{jh}(l)\Sigma^2_{jh}(l)}-1+O_p\left(\frac{1}{n_{jh}}\right)\right]\\
&= \frac{\sum_{i:c_i=h,x_{ij}=l}\{y_i-M_{jh}(l)\}^2}{2\Sigma^2_{jh}(l)}-\frac{n_{jh}(l)}{2}+O_p(1) \qquad (16)
\end{aligned}$$

We now consider the case H_{0j}: $\Omega_{jh}(l)=\Omega_h, M_{jh}(l)=\mu_h, \Sigma_{jh}(l)=\sigma_h$. Since $n_{jh}(l)=p_{jhl}n^l_j$, we know by (13),

$$\begin{aligned}
\sum_{l=0}^{1}\log\left[\left(\frac{p_{jhl}}{\omega_h}\right)^{n_{jh}(l)}\right] &= \sum_{l=0}^{1}n^l_j p_{jhl}\log\left(\frac{p_{jhl}}{\omega_h}\right)\\
&= \sum_{l=0}^{1}n^l_j\omega_h\left\{\frac{p_{jhl}}{\omega_h}-1+O_p\left(\frac{1}{n^l_j}\right)\right\}\\
&= n_{jh}(0)+n_{jh}(1)-n\omega_h+O_p(1).
\end{aligned}$$

It can be easily verified that

$$\sum_{h=1}^{k}\sum_{l=0}^{1}[n_{jh}(0)+n_{jh}(1)-n\omega_h]=0.$$

Hence

$$\sum_{h=1}^{k}\sum_{l=0}^{1}\log\left[\left(\frac{p_{jhl}}{\omega_h}\right)^{n_{jh}(l)}\right]=O_p(1). \qquad (17)$$

This implies

$$\begin{aligned}\log(\mathrm{BF}_{11}(j)) &\sim -(k-1)\log(n) + \sum_{h=1}^{k}\sum_{l=0}^{1}\log\left[\left(\frac{p_{jhl}}{\omega_h}\right)^{n_{jh}(l)}\right] \\ &= -(k-1)\log(n) + O_p(1).\end{aligned}$$

Thus $\mathrm{BF}_{11}(j) \to 0$ with rate $n^{-(k-1)}$.

On the other hand, by (16),

$$\begin{aligned}\log[\mathrm{BF}_{12}(j)] &= -\frac{n}{2} - 2k\log(n) + \sum_{h=1}^{k}\sum_{l=0}^{1}\log\left[\left\{\frac{n_{jh}(l)}{2b_{jh}(l)}\right\}^{\frac{n_{jh}(l)}{2}}(\sigma_h^2)^{\frac{n_{jh}(l)}{2}}\right] \\ &\quad + \sum_{h=1}^{k}\sum_{i:c_i=h}\frac{(y_i-\mu_h)^2}{2\sigma_h^2} + O_p(1) \\ &= -\frac{n}{2} - 2k\log(n) + \sum_{h=1}^{k}\sum_{l=0}^{1}\left\{-\frac{\sum_{i:c_i=h,x_{ij}=l}\{y_i-\mu_h\}^2}{2\sigma_\mu^2} + \frac{n_{jh}(l)}{2}\right\} \\ &\quad + \sum_{h=1}^{k}\sum_{i:c_i=h}\frac{(y_i-\mu_h)^2}{2\sigma_h^2} + O_p(1) \\ &= -2k\log(n) + O_p(1).\end{aligned}$$

Thus $\mathrm{BF}_{11}(j) \to 0$ with rate n^{-2k}. Combining BF_{11} and BF_{12}, then $\Pr(H_{1j}|y,\Phi) \to 0$ with rate $n^{-(k-1)}$ under H_{0j}.

Now consider the case H_{1j}. We know by (17)

$$\begin{aligned}\log(\mathrm{BF}_{11}(j)) &\sim -(k-1)\log(n) + \sum_{h=1}^{k}\sum_{l=0}^{1} n_{jh}(l)\log\left(\frac{p_{jhl}}{\omega_h}\right) \\ &= -(k-1)\log(n) + \sum_{h=1}^{k}\sum_{l=0}^{1} n_{jh}(l)\left\{\log\frac{\Omega_{jh}(l)}{\omega_h} + \log\frac{p_{jhl}}{\Omega_{jh}(l)}\right\} \\ &= -(k-1)\log(n) + nA + O_p(1),\end{aligned}$$

where

$$A = \sum_{h=1}^{k}\left[\lambda_0\Omega_{jh}(0)\log\left(\frac{\Omega_{jh}(0)}{\omega_h}\right) + \lambda_1\Omega_{jh}(1)\log\left(\frac{\Omega_{jh}(1))}{\omega_h}\right)\right].$$

The sum A can be seen as Kullback-Leibler divergence between distributions $P = \{\lambda_0\Omega_{jh}(0), \lambda_1\Omega_{jh}(1)\}_{h=1}^{k}$

and $Q = \{\lambda_0\omega_h, \lambda_1\omega_h\}_{h=1}^k$. Under the condition $\Omega_{jh}(0) \neq \Omega_{jh}(1)$ for some h, the Kullback-Leibler divergence $A > 0$. Therefore $\mathrm{BF}_{11}(j)$ diverges to ∞ with rate $e^{nA}n^{-(k-1)}$.

Next for $\log(\mathrm{BF}_{12}(j))$, we know by (16) and (17) that

$$\begin{aligned}
\log(\mathrm{BF}_{12}(j)) &= -\frac{n}{2} - 2k\log(n) + \sum_{h=1}^{k}\sum_{l=0}^{1}\log\left[\left\{\frac{n_{jh}(l)}{2b_{jh}(l)}\right\}^{\frac{n_{jh}(l)}{2}}(\Sigma_{jh}^2(l))^{\frac{n_{jh}(l)}{2}}\right] \\
&\quad + \sum_{h=1}^{k}\sum_{l=0}^{1}\frac{n_{jh}(l)}{2}\log\left[\frac{\sigma_h^2}{\Sigma_{jh}^2(l)}\right] + \sum_{h=1}^{k}\sum_{i:c_i=h}\frac{(y_i-\mu_h)^2}{2\sigma_h^2} + O_p(1) \\
&= -\frac{n}{2} - 2k\log(n) + \sum_{h=1}^{k}\sum_{l=0}^{1}\frac{n_{jh}(l)}{2}\log\left[\frac{\sigma_h^2}{\Sigma_{jh}^2(l)}\right] \\
&\quad + \sum_{h=1}^{k}\sum_{i:c_i=h}\frac{(y_i-\mu_h)^2}{2\sigma_h^2} + O_p(1).
\end{aligned}$$

We decompose the sum

$$\begin{aligned}
&\sum_{h=1}^{k}\sum_{i:c_i=h}\frac{(y_i-\mu_h)^2}{\sigma_h^2} \\
&= \sum_{h=1}^{k}\sum_{l=0}^{1}\sum_{i:c_i=h,x_{ij}=l}\left[\frac{\{y_i - M_{jh}(l)\}^2 + 2\{y_i - M_{jh}(l)\}\{M_{jh}(l)-\mu_h\} + \{M_{jh}(l)-\mu_h\}^2}{\sigma_h^2}\right] \\
&= \sum_{h=1}^{k}\sum_{l=0}^{1}\left[\frac{n_{jh}(l)\{M_{jh}(l)-\mu_h\}^2 + 2n_{jh}(l)\{\bar{y}_{jh}(l) - M_{jh}(l)\}\{M_{jh}(l)-\mu_h\}}{\sigma_h^2}\right] \\
&\quad + \sum_{h=1}^{k}\sum_{l=0}^{1}\left[\frac{\Sigma_{jh}^2(l)}{\sigma_h^2}\sum_{i:c_i=h,x_{ij}=l}\frac{\{y_i - M_{jh}(l)\}^2}{\Sigma_{jh}^2(l)}\right] \\
&= \sum_{h=1}^{k}\sum_{l=0}^{1}\left[\frac{n\lambda_l p_{jhl}\{M_{jh}(l)-\mu_h\}^2}{\sigma_h^2} + O_p(\sqrt{n}) + \frac{\Sigma_{jh}^2(l)}{\sigma_h^2}\left\{n\lambda_l p_{jhl} + O_p(\sqrt{n})\right\}\right] \\
&= n\sum_{h=1}^{k}\sum_{l=0}^{1}\frac{\lambda_l\Omega_{jh}(l)\{M_{jh}(l)-\mu_h\}^2}{\sigma_h^2} + n\sum_{h=1}^{k}\sum_{l=0}^{1}\frac{\Sigma_{jh}^2(l)}{\sigma_h^2}\lambda_l\Omega_{jh}(l) + O_p(\sqrt{n}).
\end{aligned}$$

Therefore,

$$\begin{aligned}\log(\mathrm{BF}_{12}) &\sim -2k\log(n) - \frac{n}{2} + n\sum_{h=1}^{k}\sum_{l=0}^{1}\lambda_l\frac{\Omega_{jh}(l)}{2}\log\left[\frac{\sigma_h^2}{\Sigma_{jh}^2(l)}\right] \\ &\quad + \frac{n}{2}\sum_{h=1}^{k}\sum_{l=0}^{1}\frac{\lambda_l\Omega_{jh}(l)\{M_{jh}(l)-\mu_h\}^2}{\sigma_h^2} + \frac{n}{2}\sum_{h=1}^{k}\sum_{l=0}^{1}\frac{\Sigma_{jh}^2(l)}{\sigma_h^2}\lambda_l\Omega_{jh}(l) + O_p(\sqrt{n}) \\ &= nB + O_p(\sqrt{n}),\end{aligned}$$

where

$$B = \frac{1}{2}\sum_{h=1}^{k}\sum_{l=0}^{1}\frac{\lambda_l\Omega_{jh}(l)\{M_{jh}(l)-\mu_h\}^2}{\sigma_h^2} + \frac{1}{2}\sum_{h=1}^{k}\sum_{l=0}^{1}\lambda_l\Omega_{jh}(l)\left[\frac{\Sigma_{jh}^2(l)}{\sigma_h^2} - 1 + \log\left\{\frac{\sigma_h^2}{\Sigma_{jh}^2(l)}\right\}\right].$$

We see for $t > 0$,

$$t - 1 - \log t \geq 0,$$

as $t - 1 - \log t$ has minimum value 0 at $t = 1$. This means that

$$\frac{\Sigma_{jh}^2(l)}{\sigma_h^2} - 1 + \log\left[\frac{\sigma_h^2}{\Sigma_{jh}^2(l)}\right] \geq 0,$$

with equality holding if and only if $\Sigma_{jh}^2(l) = \sigma_h^2$. Therefore, $B \geq 0$ and $B = 0$ if and only if $M_{jh}(0) = M_{jh}(1) = \mu_h, \Sigma_{jh}^2(0) = \Sigma_{jh}^2(1) = \sigma_h^2$ for all $h = 1, 2, ...k$. This implies under the condition $\theta_j(0) \neq \theta_j(1)$, $\log \mathrm{BF}_{12}(j) \to \infty$ with rate of Bn. Under H_1, then either $\omega_j(0) \neq \omega_j(1)$ or $\theta_j(0) \neq \theta_j(1)$. In either case, at least one of $\log \mathrm{BF}_{11}(j)$ or $\log \mathrm{BF}_{12}(j)$ goes to infinity with rate An or Bn and hence $\Pr(H_{0j}|y, x_j, \Lambda, \Phi) \to 0$ with rate of $e^{-n\max(A,B)}$. □

Proof of Lemma 3.2

Lemma 6.2 . *Let $y_1, \ldots, y_n$ be independent with density f_0. Assume $f^* = argmin_{f\in\mathbb{F}} KL(f_0||f)$ exists and $P_n(B(\epsilon, f^*|f_0)) > 0$ for all $\epsilon > 0$. Then, for any fixed $\epsilon > 0$,*

$$P_n(f \in \mathbb{F} : d(f, f^*) \geq \epsilon|y) \to 0.$$

Proof. The proof follows from Theorem 2 from Lock and Dunson (2015). Note that the positive

neighborhood condition can be satisfied since ω has a Dirichlet prior, which induces positive support over $\mathbb{F}$. □

The uniqueness assumption $d(\sum \omega_k f_k(\cdot|\theta_k), f^*) = 0$ implies $f = f^*$ and $(\omega, \theta) = (\omega^*, \theta^*)$. Because $d(\sum \omega_k f_k(\cdot|\theta_k), f^*)$ is continuous with respect to ω, θ, there exists $\delta > 0$ such that $d(\sum \omega_k f_k(\cdot|\theta_k), f^*) < \delta$ implies $||(\omega, \theta) - (\omega^*, \theta^*)|| < \epsilon$. Hence by the previous lemma:

$$\begin{aligned} P_n((\omega, \theta) \in (\mathbb{S}^{k-1}, \mathbb{R}^k) : ||(\omega, \theta) - (\omega^*, \theta^*)|| \geq \epsilon | y) &\leq P_n(f \in \mathbb{F} : d(f, f^*) \geq \epsilon | y) \\ &\to 0 \end{aligned} \tag{18}$$

. □

Proof of Theorem 3.3

Since $\Pr(H_{0j}|y, x_j, \omega, \theta, \Lambda)$ is a continuous function of ω and θ, and $(\omega, \theta) \to (\omega^*, \theta^*)$ by the previous theorem, we know that $\Pr(H_{0j}|y, x_j, \omega, \theta, \Lambda) - \Pr(H_{0j}|y, x_j, \omega^*, \theta^*, \Lambda) \to 0$ when $n_j^0, n_j^1 \to \infty$.

Under H_{0j}, $f^{(0)} = f^{(1)}$, so it follows that $f^{*(0)} = f^{*(1)} = f^*$, and $\omega^{*(0)} = \omega^{*(1)} = \omega^*$, $\theta^{*(0)} = \theta^{*(1)} = \theta^*$. By Theorem 3.1, $\Pr(H_{0j}|y, x_j, \omega^*, \theta^*, \Lambda) \to 1$, thus $\Pr(H_{0j}|y, x_j, \omega, \theta, \Lambda) \to 1$ and hence $\Pr(H_{0j}|y, x_j, \Lambda) \to 1$.

On the other hand, if $f^{*(0)} \neq f^{*(1)}$ under H_{1j}, then by Theorem 3.1, $\Pr(H_{1j}|y, x_j, \omega^*, \theta^*, \Lambda) \to 0$, and thus $\Pr(H_{0j}|y, x_j, \omega, \theta, \Lambda) \to 0$ and $\Pr(H_{0j}|y, x_j, \Lambda) \to 0$. □

Proof of Theorem 3.4

By (C1), all jointly important predictors are marginally important which implies that the alternative hypothesis H_{1j} is true for all $j \in D$. Hence by (C2) and Theorem 3.3, $\Pr(H_{0j}|y, x_j, \Lambda) = 0$. Thus asymptotically as $n, p \to \infty$, $P(D \subset \hat{D}) \to 1$. □